

Paul He

Toronto, ON | hepaul@cs.toronto.edu | +1 416 659 5646 | hepaul7.github.io | github.com/hepaul7

Education

University of Toronto , BSc in Computer Science (Honours)	Sep 2021 – May 2025
• Graduating with High Distinction	
ETH Zürich , Exchange Year in Computer Science	Sep 2023 – Aug 2024
• Completed graduate-level coursework in NLP and Machine Learning	

Research Experience

Sensitive Functions and the Limits of Self-Attention in NLP	Jan 2025 – Present
<i>Supervisor: Prof. Gerald Penn</i>	University of Toronto
• Studying the presence and impact of sensitive Boolean functions (e.g., parity) in natural language understanding tasks.	
• Analyzing whether models like BERT can represent such functions or fail due to attention limitations.	
Causal Reasoning Evaluation for Language Models	Jan 2025 – Present
<i>Supervisor: Prof. Zhijing Jin (Jinesis AI Lab)</i>	Vector Institute
• Developing a tool to parse and evaluate causal expressions generated by LLMs, using BFS to match equivalent causal formulas and identify reasoning errors.	
• Aiming to use this as a proxy to improve QA tasks involving causal understanding.	
• Planning to submit this work to EMNLP 2025.	
Polynomial-Time CCG Parsing: An Empirical Evaluation of the Kuhlmann–Satta Algorithm	Sep 2025 - Dec 2025
<i>Supervisor: Prof. Dr. Gerald Penn</i>	University of Toronto
• Designed and implemented a polynomial-time parser for Combinatory Categorical Grammar (CCG).	
• Conducted large-scale multilingual experiments to analyze when symbolic parsers outperform naive baselines, identifying empirical complexity thresholds.	
• First-author paper under review at the 18th International Conference on Parsing Technologies (IWPT 2025).	
Unsupervised Prompt Optimization for RAG using PMI	Jun 2024 – Sep 2024
<i>Supervisors: Tianyu Liu, Prof. Ryan Cotterell (Rycolab)</i>	ETH Zürich
• Investigated attribution methods such as change in Jensen-Shannon divergence for Fisher information to approximate attention weights in LLMs.	
• Developed and implemented an unsupervised method to reorder retrieved documents using PMI, improving QA accuracy by up to 4% on NQ-Open and ELI5.	
• Presented the work at NAACL 2025 poster session, attracting sustained interest from researchers across academia and industry.	

Publications

Pointwise Mutual Information as a Performance Gauge for Retrieval-Augmented Generation	2025
Tianyu Liu, Jirui Qi, Paul He , Arianna Bisazza, Mrinmaya Sachan, Ryan Cotterell	
NAACL 2025 (Main), https://aclanthology.org/2025.naacl-long.78/	

Projects

Speeding up Earley’s Algorithm	Mar 2024 - Aug 2024
• Proposed and implemented a grid search approach, which resulted in an algorithm $\times 5$ faster than the naive algorithm.	
• Tested and experimented with the proposed algorithms on Treebank datasets to validate theoretical findings.	
Word Segmentation and Part-of-Speech Tagging with Transformer	Oct 2023 - Jan 2024
• Designed a new decoder embedding mechanism to better capture subtle linguistic nuances and dependencies.	
• Implemented the mechanism and achieved state-of-the-art results with a validation accuracy of 96.29%.	

Skills

Programming Languages: Python, C, Java, SQL

Spoken Languages: English (Fluent), Mandarin (Proficient), German (Proficient)